From:	Dan Bornstein.	Sent:4/28/2010 11:18 AM.
To: [-]	Bill Buzbee.	
Cc: [-]	dalvik-team.	
Bcc: [-]	.	
Subject:	Re: [dalvik-team] This just in re Parfait.	

On Wed, Apr 28, 2010 at 10:35 AM, Bill Buzbee <buzbee@google.com> wrote:
> Regular comparisons can be good - let each team know if there's an
> optimization it's missing that the other team is doing. However, anything
> we set up should also include memory consumption for both VMs. This would
> involve not only the size of the JIT'd code, but the consumption of the JIT
> during compilation.

Agreed. Here's how I replied to that mail (just after forwarding it to you al
> ...Bill
>
> On Wed, Apr 28, 2010 at 10:00 AM, Dan Bornstein <danfuzz@google.com> wrote:
>>
>> It took a week of pings and re-requests (and, fwiw, see
>>
>> <https://groups.google.com/a/google.com/group/v8-team/browse_thread/thread/7c6a7694d858f6bc/2ff8f06cd9cebe5b?lnk=st&q=lars+bak#2ff8f06cd9cebe5b>
>> in re his radio silence), but this looks like progress.
>>
>> Ben, if you haven't already heard from Regis, then now is the time to
>> try to make contact again.
>>
>> -dan
>>
>> ---------- Forwarded message ----------
>> From: Lars Bak <bak@google.com>
>> Date: Wed, Apr 28, 2010 at 3:22 AM
>> Subject: Re: perf comparison?
>> To: Dan Bornstein <danfuzz@google.com>
>> Cc: Hiroshi Lockheimer <hiroshi@google.com>
>>
>>
>> Hi Dan & Hiroshi
>>
>> Sorry for the delay answering this, but I got stuck in ashes and now
>> on vacation.
>> The best way to resolve these discrepancies quickly is by setting up a
>> meeting with Regis and Srdjan in 1950 and then go through the exercise
>> of producing the numbers on our Tegra board. I assume, resolving this
>> over email will take too long.
>> I have asked Regis for the Tegra board specs and I'll forward it as
>> soon as I get it.
>> First of all we have to make sure we run the same versions of the
>> benchmarks.
>>
>> Moving forward I would like to track performance consistently between
>> the two JVM by using a tool like http://go/lem.
>> Both on ARM and Atom. I expect this will benefit both teams.
>>
>> Regards,
>> Lars
>>
>>
>> On Tue, Apr 20, 2010 at 8:25 AM, Dan Bornstein <danfuzz@google.com> wrote:
>> > On Mon, Apr 19, 2010 at 10:22 PM, Hiroshi Lockheimer
>> > <hiroshi@google.com> wrote:

>> >> Hi Lars, thanks for the info. Do you know which Tegra board was used?
>> >> Dan,
>> >> are you able to repro all this?
>> >
>> > We haven't tried to build Parfait, but we have tried to reproduce the
>> > Dalvik benchmark numbers that were reported. I'm afraid the numbers we
>> > gathered don't seem to match up with the presentation at the all-hands
>> > nor with the numbers Lars just shared, but there are at least a couple
>> > open questions the answers to which could help resolve the
>> > discrepancies.
>> >
>> > One open question is, as you mention above, what hardware the Parfait
>> > team used for their tests, or what its general specs are. The Tegra
>> > board we have for testing is likely a different make than the one the
>> > Parfait team uses, but it would be good to know for sure. Ben asked
>> > Regis to run the native (compiled C) version of DeltaBlue on his
>> > devboard to help us get a baseline for comparison, but I don't believe
>> > Ben has yet heard back from him.
>> >
>> > As an example of a discrepancy in the numbers, Lars, the DeltaBlue
>> > number you report for "Dalvik noJIT on Tegra" is 648.63 msec, but the
>> > number we measured for our Tegra board for the Dalvik interpreter was
>> > 35.47 msec. And our actual measurement of DeltaBlue on a Tegra build
>> > of Dalvik with JIT was 18.85 msec, whereas the Parfait Tegra
>> > measurement you report was 121.93 msec. My expectation is that Parfait
>> > would actually be faster than Dalvik, but instead based on this number
>> > it looks like Dalvik is 6.5x Parfait, at least for this benchmark.
>> > This all makes me think we aren't actually measuring the same thing,
>> > or even anything close to the same thing.
>> >
>> > Other than hardware, here are some quick thoughts on possible
>> > differences there might be between our respective environments:
>> > Different versions of the benchmark code (maybe you can share the
>> > sources or .jars of what you're running), how the code was processed
>> > into .dex files for the purposes of testing Dalvik (dx has options
>> > that can affect code quality), and how Dalvik itself was configured
>> > (e.g. was it in a position to do install-time optimization?). I would
>> > be happy to furnish the Parfait team with our versions of any of this.
>> >
>> > Thanks in advance for the help.
>> >
>> > -dan
>> >
>
>